arXiv:1604.04280v2 [q-bio.GN] 22 Apr 2016

VARIATIONAL INFERENCE FOR RARE VARIANT DETECTION IN DEEP, HETEROGENEOUS NEXT-GENERATION SEQUENCING DATA

FAN ZHANG [1], PATRICK FLAHERTY [1,2]

[1] *Department of Biomedical Engineering, Worcester Polytechnic Institute, MA, USA*
[2] *Department of Mathematics and Statistics, University of Massachusetts, Amherst, MA, USA*

ABSTRACT. The detection of rare variants is important for understanding the genetic heterogeneity in mixed samples. Recently, next-generation sequencing (NGS) technologies have enabled the identification of single nucleotide variants (SNVs) in mixed samples with high resolution. Yet, the noise inherent in the biological processes involved in next-generation sequencing necessitates the use of statistical methods to identify true rare variants. We propose a novel Bayesian statistical model and a variational expectation-maximization (EM) algorithm to estimate non-reference allele frequency (NRAF) and identify SNVs in heterogeneous cell populations. We demonstrate that our variational EM algorithm has comparable sensitivity and specificity compared with a Markov Chain Monte Carlo (MCMC) sampling inference algorithm, and is more computationally efficient on tests of low coverage ($27\times$ and $298\times$) data. Furthermore, we show that our model with a variational EM inference algorithm has higher specificity than many state-of-the-art algorithms. In an analysis of a directed evolution longitudinal yeast data set, we are able to identify a time-series trend in non-reference allele frequency and detect novel variants that have not yet been reported. Our model also detects the emergence of a beneficial variant earlier than was previously shown, and a pair of concomitant variants.

1. INTRODUCTION

Massively parallel sequencing data generated by next-generation sequencing technologies is routinely used to interrogate SNVs in research samples [Koboldt *et al.*, 2013]. For example, deep sequencing confirmed the degree of genetic heterogeneity of HIV and influenza [Flaherty *et al.*, 2011; Ghedin *et al.*, 2011]. Many solid tumors are represented as having intra-tumor heterogeneity by DNA sequencing technology, where SNVs are rare in the population [Navin *et al.*, 2010]. Also, whole genome sequencing has revealed that many beneficial mutations of minor allele frequencies are essential to respond to dynamic environments [Kvitek and Sherlock, 2013]. However, rare SNVs identification in heterogeneous cell populations is challenging, because of the intrinsic error rate of next generation sequencing platform [Shendure and Ji, 2008]. Thus, there is a need for accurate and scalable statistical methods to uncover SNVs in heterogeneous samples.

A number of computational methods have been developed to detect SNVs in large scale genomic data sets. These methods can be roughly categorized as probabilistic or heuristic or some combination. Among all of the current probabilistic methods, the Bayesian probabilistic framework has been increasingly used to calculate the unobserved quantities such as variant allele frequency given the observed genomic sequencing data. GATK [McKenna *et al.*, 2010] and SAMTools [Li *et al.*, 2009] use a naive Bayesian decision rule to call variants. EBCall models sequencing errors based on a Beta-Binomial distribution, where the parameters and latent variables are estimated from a set of non-paired normal sequencing samples [Shiraishi *et al.*, 2013]. However, the error rate of normal sequencing samples could be unmatched with the error rate of the target samples, which may cause a problem of making false negatives calls [Wang *et al.*, 2013]. CRISP compares aligned reads across multiple pools to obtain sequencing errors, and then distinguishes true rare variants

from the sequencing errors [Bansal, 2010]. However, the bottleneck of CRISP is its low computational efficiency due to a calculation of a large amount of contingency tables.

Generally, independent modeling method models the genotypes of tumour and normal samples separately and then looks for the difference between them. In contrast to independently analyzing a tumour-normal pair, joint modeling method models a joint-genotype of the tumour and normal samples simultaneously. JointSNVMix introduces two Bayesian probabilistic models (JointSNVMix1 and JointSNVMix2) to jointly analyze a tumour-normal paired allelic count of NGS data [Roth *et al.*, 2012]. JointSNVMix derives an expectation maximization (EM) algorithm to calculate maximum a-posteriori (MAP) estimate of latent variables in a particular probabilistic graphical model. Furthermore, Roth *et al.* [2012] reveals that the joint modeling method, JointSNVMix1, observes 80-fold reduction of false positives compared with its independent analogue (SNVMix1). SomaticSniper [Larson *et al.*, 2012] models the joint diploid genotype likelihoods for both tumour and normal samples. Additionally, Strelka [Saunders *et al.*, 2012] models the joint probabilistic distribution of allele frequencies for both tumour and normal samples, which is demonstrated to be more accurate compared with the methods based on the estimated allele frequency tests between tumour and normal samples.

An alternative strategy to probabilistic methods is heuristic methods that use a set of criteria to select variant positions instead of modeling the data using probabilistic distributions. VarScan is an example of a heuristic method that compares tumour and normal samples by satisfying some lower bounds, such as a certain variant allele frequency and a number of allele counts [Koboldt *et al.*, 2012]. SNVer focuses on a frequentist method that is able to calculate P-values, but Wei *et al.* [2011] pointed out that this approach fails to model sampling bias that will reduce the power of detecting true rare variants.

In previous work, we developed a Beta-Binomial model to characterize a null hypothesis error rate distribution at each position. Using this rare variant detection (RVD) model, we call rare variants by comparing the error rate of the sample sequence data to a null distribution obtained from sequencing a known reference sample [Flaherty *et al.*, 2011]. RVD can identify mutant positions at a 0.1% fraction in mixed samples using high read depth data.

We improved upon that work, RVD2, by using hierarchical priors to tie parameters across positions [He *et al.*, 2015]. We derived a Markov Chain Monte Carlo (MCMC) sampling algorithm for posterior inference. However, the main limitation of MCMC is that it is hard to diagnose convergence and may be slow to converge [Jordan *et al.*, 1999]. An alternative method, that we explore here, is to use variational inference, which is based on a proposed variational distribution over latent variables. By optimizing variational parameters, we fit an approximate distribution that is close to the true posterior distribution in the sense of the Kullback-Liebler (KL) divergence. Variational inference can now handle nonconjugate distributions and tends to be more computationally efficient than MCMC sampling [Peterson and Hartman, 1989].

Here, we propose a variational EM algorithm for our Bayesian statistical model to detect SNVs in heterogeneous NGS data. We show that variational EM algorithm has comparable accuracy and efficiency compared with MCMC in a synthetic data set. In section 2, we define the model structure. In section 3, we derive our variational EM algorithm to approximate the posterior distribution over latent variables. In section 4, we call a variant by a posterior difference hypothesis test between the key model parameters of a pair of samples. In section 5, we compare the performance of the variational EM inference algorithm to the MCMC sampling method and the state-of-the-art methods using a synthetic data set. We also show that our variational EM algorithm is able to detect rare variants and estimate non-reference allele frequency (NRAF) in a longitudinal directed evolution experimental data set.

2. Model Structure

Our Bayesian statistical model is shown as a graphical model in Figure 1. In the model, r_{ji} is the number of reads with a non-reference base at location j in experimental replicate i; n_{ji} is the total number of reads at location j in experimental replicate i. The model parameters are:

μ_0**:** a global non-reference read rate that captures the error rate across all the positions,
M_0**:** a global precision that captures the variation of the error rate across positions in a sequence,
M_j**:** a local precision that captures the variation of the error rate at position j across different replicates.

The latent variables are:

$\mu_j \sim$ **Beta**(μ_0, M_0)**:** position-specific non-reference read rate for position j,
$\theta_{ji} \sim$ **Beta**(μ_j, M_j)**:** the non-reference read rate for position j in replicate i.

In Figure 1B, γ is the parameter for the variational distribution for latent variable μ, and δ is the parameter for the variational distribution for latent variable θ. We describe $q(\mu)$ and $q(\theta)$ in detail in section 3.3

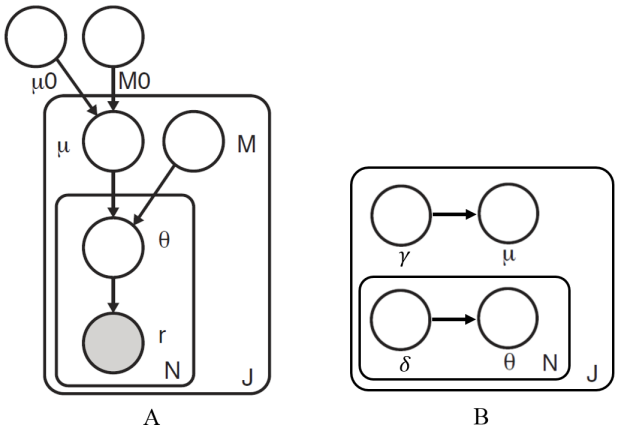

FIGURE 1. A. Graphical model representation of the model. B. Graphical model representation of the variational approximation to approximate the posterior distribution. Observed random variables are shown as shaded nodes and latent random variables are unshaded. The object of inference for the variational EM algorithm is the joint distribution $p(\mu, \theta | r, n)$.

The model generative process is as follows:

(1) For each location $j \in [1, \ldots, J]$:
  (a) Draw an error rate $\mu_j \sim$ Beta(μ_0, M_0)
  (b) For each replicate $i \in [1, \ldots, N]$:
    (i) Draw $\theta_{ji} \sim$ Beta(μ_j, M_j)
    (ii) Draw $r_{ji} | n_{ji} \sim$ Binomial(θ_{ji}, n_{ji})

The joint distribution $p(r, \mu, \theta | n; \phi)$ given the parameters can be factorized as

$$p(r, \mu, \theta | n; \phi) = p(r | \theta, n) p(\theta | \mu; M) p(\mu; \mu_0, M_0). \tag{1}$$

3. Variational Expectation Maximization (EM) Inference

We developed a non-conjugate variational inference algorithm to approximate the posterior distribution,

$$p(\mu, \theta | r, n; \phi) = \frac{p(r, \mu, \theta | n; \phi)}{p(r | n; \phi)}, \tag{2}$$

where the parameters are $\phi \triangleq \{\mu_0, M_0, M\}$.

3.1. **Factorization.** We propose the following factorized variational distribution to approximate the true posterior over latent variables μ_j and θ_{ji}. Here, $q(\mu_j)$ approximates the variational posterior distribution of μ_j, which represents the local error rate distribution in position j across different replicates; and $q(\theta_{ji})$ approximates the posterior distribution of θ_{ji}, which is the error rate distribution in position j replicate i.

$$q(\mu, \theta) = q(\mu)q(\theta) = \prod_{j=1}^{J} q(\mu_j) \prod_{i=1}^{N} q(\theta_{ji}). \tag{3}$$

3.2. **Evidence Lower Bound (ELBO).** Given the variational distribution, q, the log-likelihood of the data is lower-bounded according to Jensen's inequality,

$$\begin{aligned}
\log p\left(r|\phi\right) &= \log \int_{\mu} \int_{\theta} p\left(r, \mu, \theta|n; \phi\right) d\theta d\mu \\
&= \log \int_{\mu} \int_{\theta} p\left(r, \mu, \theta|n; \phi\right) \frac{q\left(\mu, \theta\right)}{q\left(\mu, \theta\right)} d\theta d\mu \\
&\geq \int_{\mu} \int_{\theta} q\left(\mu, \theta\right) \log \frac{p\left(r, \mu, \theta|n; \phi\right)}{q\left(\mu, \theta\right)} d\theta d\mu \\
&= E_q\left[\log p\left(r, \mu, \theta|n; \phi\right)\right] - E_q\left[\log q\left(\mu, \theta\right)\right] \\
&\triangleq \mathcal{L}(q, \phi).
\end{aligned} \tag{4}$$

The function $\mathcal{L}(q, \phi)$ is the evidence of lower bound (ELBO) of the log-likelihood of the data, which is the sum of q-expected complete log-likelihood and the entropy of the variational distribution q. The goal of variational inference is to maximize the ELBO. Equivalently, q is chosen by minimizing the KL divergence between the variational distribution and the true posterior distribution.

The ELBO can be expanded as

$$\begin{aligned}
\mathcal{L}(q, \phi) &= E_q\left[\log p\left(r, \mu, \theta|n; \phi\right)\right] - E_q\left[\log q\left(\mu, \theta\right)\right] \\
&= E_q\left[\log p\left(r|\theta, n\right)\right] + E_q\left[\log p\left(\theta|\mu; M\right)\right] + E_q\left[\log p\left(\mu; \mu_0, M_0\right)\right] - E_q\left[\log q\left(\mu\right)\right] - E_q\left[\log q\left(\theta\right)\right].
\end{aligned} \tag{5}$$

We write out each component below. The detail of the derivation of $E_q\left[\log p\left(r|\theta, n\right)\right]$, $E_q\left[\log p\left(\mu; \mu_0, M_0\right)\right]$, and $E_q\left[\log p\left(\theta|\mu; M\right)\right]$ is shown in Appendix A.

$$\begin{aligned}
E_q\left[\log p\left(r|\theta, n\right)\right] &= \sum_{j=1}^{J} \sum_{i=1}^{N} E_q\left[\log p\left(r_{ji}|\theta_{ji}, n_{ji}\right)\right] \\
&= \sum_{j=1}^{J} \sum_{i=1}^{N} \log\left(\frac{\Gamma(n_{ji}+1)}{\Gamma(r_{ji}+1)\Gamma(n_{ji}-r_{ji}+1)}\right) \\
&\quad + \sum_{j=1}^{J} \sum_{i=1}^{N} \left\{r_{ji} E_q\left[\log \theta_{ji}\right] + (n_{ji}-r_{ji}) E_q\left[\log(1-\theta_{ji})\right]\right\}
\end{aligned} \tag{6}$$

$$
\begin{aligned}
E_q\left[\log p\left(\mu;\mu_0,M_0\right)\right] &= \sum_{j=1}^{J} E_q\left[\log p\left(\mu_j;\mu_0,M_0\right)\right] \\
&= J * \log\frac{\Gamma(M_0)}{\Gamma(\mu_0 M_0)\Gamma(M_0(1-\mu_0))} \\
&\quad + \sum_{j=1}^{J}\left\{(M_0\mu_0-1)E_q\left[\log\mu_j\right] + (M_0(1-\mu_0)-1)E_q\left[\log(1-\mu_j)\right]\right\}
\end{aligned} \tag{7}
$$

$$
\begin{aligned}
E_q\left[\log p\left(\theta|\mu;M\right)\right] &= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log p\left(\theta_{ji}|\mu_j;M_j\right)\right] \\
&= N * \sum_{j=1}^{J} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right] \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{M_j E_q\left[\mu_j\right] E_q\left[\log\theta_{ji}\right] - E_q\left[\log\theta_{ji}\right]\right\} \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{\left(M_j - 1 - M_j E_q\left[\mu_j\right]\right) E_q\left[\log\left(1-\theta_{ji}\right)\right]\right\}
\end{aligned} \tag{8}
$$

Therefore, we need to compute the following expectations with respect to the variational distribution: $E_q\left[\log\theta_{ji}\right]$, $E_q\left[\log\left(1-\theta_{ji}\right)\right]$, $E_q\left[\log\mu_j\right]$, $E_q\left[\log(1-\mu_j)\right]$, $E_q\left[\mu_j\right]$, and $E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right]$.

We select the functional forms for the variational distributions $q(\theta)$ and $q(\mu)$ to facilitate these expected value computations.

3.3. **Variational Distributions.** Since θ and r are conjugate pairs, the posterior distribution of θ_{ji} is a Beta distribution,

$$
p(\theta_{ji}|r_{ji},n_{ji},\mu_j,M_j) \sim \text{Beta}(r_{ji}+M_j\mu_j, n_{ji}-r_{ji}+M_j(1-\mu_j)). \tag{9}
$$

Therefore, we propose a Beta distribution with parameter vector δ_{ji} as variational distribution,

$$
\theta_{ji} \sim \text{Beta}(\delta_{ji1},\delta_{ji2}).
$$

The posterior distribution of μ_j is given by its Markov blanket,

$$
p(\mu_j|\theta_{ji},M_j,\mu_0,M_0) \propto p(\mu_j|\mu_0,M_0)p(\theta_{ji}|\mu_j,M_j). \tag{10}
$$

This is not in the form of any known distribution. But, since the support of μ_j is $[0,1]$, we propose a Beta distribution with parameter vector γ_j as variational distribution,

$$
\mu_j \sim \text{Beta}(\gamma_{j1},\gamma_{j2}).
$$

Given these variational distributions, we have

$$
\begin{aligned}
E_q\left[\log\theta_{ji}\right] &= \psi(\delta_{ji1}) - \psi(\delta_{ji1}+\delta_{ji2}) \\
E_q\left[\log\left(1-\theta_{ji}\right)\right] &= \psi(\delta_{ji2}) - \psi(\delta_{ji1}+\delta_{ji2}) \\
E_q\left[\mu_j\right] &= \frac{\gamma_{j1}}{\gamma_{j1}+\gamma_{j2}} \\
E_q\left[\log\mu_j\right] &= \psi(\gamma_{j1}) - \psi(\gamma_{j1}+\gamma_{j2}) \\
E_q\left[\log(1-\mu_j)\right] &= \psi(\gamma_{j2}) - \psi(\gamma_{j1}+\gamma_{j2}),
\end{aligned} \tag{11}
$$

where ψ is the digamma function.

Since there is no analytical representation for $E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right]$, we must resort to numerical integration,

$$E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma((1-\mu_j)M_j)}\right)\right] = \int_0^1 q(\mu_j;\gamma_{j1},\gamma_{j2})\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma((1-\mu_j)M_j)}\right)d\mu_j. \quad (12)$$

Here $q(\mu_j;\gamma_{j1},\gamma_{j2})$ is the probability density function of the Beta distribution that is calculated using the Python built-in function `scipy.stats.beta.pdf`, and $\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma((1-\mu_j)M_j)}\right)$ is calculated using the Python built-in function `scipy.special.betaln`. Unfortunately, this numerical integration step is computationally expensive. Finally, the entropy terms can be computed as follows,

$$\begin{aligned} E_q\left[\log q\left(\mu\right)\right] &= \sum_{j=1}^{J} E_q\left[\log q(\mu_j)\right] \\ &= -\sum_{j=1}^{J}\left\{\log(B(\gamma_{j1},\gamma_{j2})) - (\gamma_{j1}-1)\psi(\gamma_{j1})\right\} \\ &\quad + \sum_{j=1}^{J}\left\{-(\gamma_{j2}-1)\psi(\gamma_{j2}) + (\gamma_{j1}+\gamma_{j2}-2)\psi(\gamma_{j1}+\gamma_{j2})\right\}; \end{aligned} \quad (13)$$

and

$$\begin{aligned} E_q\left[\log q\left(\theta\right)\right] &= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log q(\theta_{ji})\right] \\ &= -\sum_{j=1}^{J}\sum_{i=1}^{N}\left\{\log(B(\delta_{ji1},\delta_{ji2})) - (\delta_{ji1}-1)\psi(\delta_{ji1})\right\} \\ &\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{-(\delta_{ji2}-1)\psi(\delta_{ji2}) + (\delta_{ji1}+\delta_{ji2}-2)\psi(\delta_{ji1}+\delta_{ji2})\right\}. \end{aligned} \quad (14)$$

3.4. **Variational EM Algorithm.** Variational EM algorithm maximizes the ELBO on the true likelihood by alternating between maximization over q (E-step) and maximization over $\phi = \{\mu_0, M_0, M\}$ (M-step).

3.4.1. *(E-step): Updating the variational distributions.* The terms in the ELBO that depend on $q(\theta_{ji}|\delta_{ji1},\delta_{ji2})$ are

$$\begin{aligned} \mathcal{L}_{[q(\theta_{ji})]} &= \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{r_{ji}E_q\left[\log\theta_{ji}\right] + (n_{ji}-r_{ji})E_q\left[\log(1-\theta_{ji})\right]\right\} \\ &\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{M_j E_q\left[\mu_j\right]E_q\left[\log\theta_{ji}\right] - E_q\left[\log\theta_{ji}\right] + \left(M_j - 1 - M_j E_q\left[\mu_j\right]\right)E_q\left[\log\left(1-\theta_{ji}\right)\right]\right\} \\ &\quad - \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log q(\theta_{ji})\right] \end{aligned} \quad (15)$$

We update the variational parameters by numerically optimizing $\hat{\delta}_{ji1}, \hat{\delta}_{ji2} = \arg\max_{\delta_{ji1},\delta_{ji2}} \mathcal{L}_{[q(\theta_{ji})]}$ subject to the constraints that $\delta_{ji1} \geq 0$ and $\delta_{ji2} \geq 0$ and conditioned on fixed values for the other model and variational parameters using Sequential Least SQuares Programming (SLSQP) [Kraft and others, 1988].

We update the variational distribution $q(\mu_j)$ using the partial ELBO depending on γ_j from each position j (16).

$$\begin{aligned}\mathcal{L}_{[q(\mu_j)]} = & N \sum_{j=1}^{J} E_q \left[\log \left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))} \right) \right] \\ & + \sum_{j=1}^{J} \sum_{i=1}^{N} \left\{ M_j E_q\left[\mu_j\right] E_q\left[\log \theta_{ji}\right] - E_q\left[\log \theta_{ji}\right] + \left(M_j - 1 - M_j E_q\left[\mu_j\right]\right) E_q\left[\log\left(1-\theta_{ji}\right)\right] \right\} \\ & + J \log \frac{\Gamma(M_0)}{\Gamma(\mu_0 M_0)\Gamma(M_0(1-\mu_0))} \\ & + \sum_{j=1}^{J} \left\{ (M_0\mu_0 - 1) E_q\left[\log \mu_j\right] + (M_0(1-\mu_0) - 1) E_q\left[\log(1-\mu_j)\right] \right\} \\ & - \sum_{j=1}^{J} E_q\left[\log q(\mu_j)\right]\end{aligned} \tag{16}$$

Again, we update the variational parameters by numerically optimizing $\hat{\gamma}_{j1}, \hat{\gamma}_{j2} = \arg\max_{\gamma_{j1},\gamma_{j2}} \mathcal{L}_{[q(\mu_j)]}$ subject to the constraints that $\gamma_{j1} \geq 0$ and $\gamma_{j2} \geq 0$ and conditioned on fixed values for the other model and variational parameters using SLSQP [Kraft and others, 1988]. The computational cost of optimizing (16) is high because of the quadrature of $E_q \left[\log \left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))} \right) \right]$ in (12).

3.4.2. *(M-step): Updating the model parameters.* We can write out the ELBO as a function of each model parameter μ_0, M_0, and M_j as follows.

The ELBO with respect to μ_0 is

$$\mathcal{L}_{[\mu_0]} = -J * \log\Gamma(\mu_0 M_0) - J * \log\Gamma(M_0(1-\mu_0)) + M_0\mu_0 \sum_{j=1}^{J} \left\{ E_q\left[\log \mu_j\right] - E_q\left[\log(1-\mu_j)\right] \right\}. \tag{17}$$

The ELBO with respect to M_0 is

$$\mathcal{L}_{[M_0]} = J * \log \frac{\Gamma(M_0)}{\Gamma(\mu_0 M_0)\Gamma(M_0(1-\mu_0))} + M_0 \sum_{j=1}^{J} \left\{ \mu_0 E_q\left[\log \mu_j\right] + (1-\mu_0) E_q\left[\log(1-\mu_j)\right] \right\}. \tag{18}$$

The ELBO with respect to M_j is

$$\begin{aligned}\mathcal{L}_{[M_j]} = & N * \sum_{j=1}^{J} E_q \left[\log \left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))} \right) \right] \\ & + M_j \sum_{j=1}^{J} \sum_{i=1}^{N} \left\{ E_q\left[\mu_j\right] E_q\left[\log \theta_{ji}\right] + \left(1 - E_q\left[\mu_j\right]\right) E_q\left[\log\left(1-\theta_{ji}\right)\right] \right\}.\end{aligned} \tag{19}$$

We also use SLSQP to optimize the ELBO function with respect to each parameter, μ_0, M_0, and M_j. It is computationally easy to optimize μ_0 (17) and M_0 (18). However, it is costly for optimizing M_j (19) because the quadrature is needed to calculate $E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right]$ using (12).

The variational EM algorithm is summarized using pseudocode in Algorithm 1.

Algorithm 1 Variational EM Inference

Initialize $q(\theta, \mu)$ and $\hat{\phi}$
// M-step
repeat
  // E-step
  repeat
    for j = 1 to J **do**
      for i = 1 to N **do**
        Optimize $\mathcal{L}(q, \hat{\phi})$ over $q(\theta_{ji}; \delta_{ji}) = \text{Beta}(\delta_{ji})$
      end for
    end for
    for j = 1 to J **do**
      Optimize $\mathcal{L}(q, \hat{\phi})$ over $q(\mu_j; \gamma_j) = \text{Beta}(\gamma_j)$
    end for
  until change in $\mathcal{L}(q, \hat{\phi})$ is small
  Set $\hat{\phi} \leftarrow \arg\max_{\phi} \mathcal{L}(\hat{q}, \phi)$
until change in $\mathcal{L}(\hat{q}, \phi)$ is small

4. Hypothesis Testing

The posterior distribution over $\mu_j^{\triangle} \mid r^{case}, r^{control} \triangleq \mu_j|r^{case} - \mu_j|r^{control}$ is the distribution over the change in the non-reference read rate at position j between a case and control sample. However, we cannot compute the posterior distribution, $\mu_j^{\triangle}|r^{case}, r^{control}$, because we do not have access to the component posterior distributions $\mu_j|r^{case}$ and $\mu_j|r^{case}$. But, by approximating the distribution of $\mu_j|r^{case}$ and $\mu_j|r^{control}$ as Gaussians with moments matched to the variational posterior distributions, the distribution of $\mu_j^{\triangle}|r^{case}, r^{control}$ can be approximated as a Gaussian.

Under the variational approximation,

$$E_q[\mu_j|r^{case}] = \frac{\gamma_{j1}^{case}}{\gamma_{j1}^{case} + \gamma_{j2}^{case}} \tag{20}$$

$$\text{Var}_q[\mu_j|r^{case}] = \frac{\gamma_{j1}^{case}\gamma_{j2}^{case}}{(\gamma_{j1}^{case} + \gamma_{j2}^{case} + 1)(\gamma_{j1}^{case} + \gamma_{j2}^{case})^2} \tag{21}$$

for $\mu_j|r^{case}$ and likewise for $\mu_j|r^{control}$. We approximate the posterior for the case sample as

$$\mu_j|r^{case} \sim \mathcal{N}(E_q[\mu_j|r^{case}], \text{Var}_q[\mu_j|r^{case}]) \tag{22}$$

and likewise for the control. Then,

$$\mu_j^{\triangle} \mid r^{case}, r^{control} \sim \mathcal{N}(E_q[\mu_j|r^{case}] - E_q[\mu_j|r^{control}], \text{Var}_q[\mu_j|r^{case}] + \text{Var}_q[\mu_j|r^{control}]) \tag{23}$$

Now, we can approximate the posterior probability of interest,

$$\Pr(\mu_j^{\triangle} \geq \tau \mid r^{case}, r^{control}), \tag{24}$$

that is, the posterior probability that the difference in the non-reference read rate is greater than a fixed effect size τ (e.g. zero) for a one sided test. For a two sided test, we compute the approximate probability

$$\Pr(|\mu_j^{\triangle}| \geq \tau \mid r^{case}, r^{control}). \tag{25}$$

A position is called a *provisional variant* if $\Pr(|\mu_j^{\triangle}| \geq \tau \mid r^{case}, r^{control}) \geq 1 - \alpha/2$, where the probability is approximated as described.

It is possible that a position is called a variant due to a differential non-reference read count, but no particular alternative base is more frequently observed than the others. In this case, the likely case is a sequencing error that indiscriminately incorporates a non-reference base at the position. To discriminate this non-biological cause from the interesting true variants we use a χ^2 goodness-of-fit test for non-uniform base distribution [Efron, 2010; He *et al.*, 2015]. For each provisional variant, if we reject the null hypothesis that the distribution is uniform, we promote the position to a called variant.

5. RESULTS

5.1. **Data Sets.**

5.1.1. *Synthetic DNA Sequence Data.* Two random 400bp DNA sequences were synthesized. One sample with 14 loci variants is taken as the case and the other without variants is taken as the control. Case and control samples were mixed to yield defined NRAF of 0.1%, 0.3%, 1.0%, 10.0%, and 100.0%. The details of the experimental protocol are available in Flaherty *et al.* [2011]. The synthetic DNA data were downsampled by $10\times$, $100\times$, $1,000\times$, and $10,000\times$ using `picard` (v 1.96). The final data set contains read pairs for six replicates for the control and cases at different NRAF levels.

5.1.2. *Longitudinal Directed Evolution Data.* The longitudinal yeast data comes from three strains of haploid S288c which were grown for 448 generations under limited-glucose (0.08%). The wild-type ancestral strain GSY1136 was sequenced as a reference. Aliquots were taken about every 70 generations and sequenced. The detail of library sequencing is described in Kvitek and Sherlock [2013], Bansal [2010], and Kao and Sherlock [2008]. The Illumina sequencing data are available on the NCBI Sequence Read Archive (SRA054922) [Kvitek and Sherlock, 2013]. For this study, we received the original BAM files from one of the authors. The aligned BAM files have 266 - $1,046\times$ coverage. We used `samtools` (v 1.1) with `-mpileup -C50` flags to convert BAM files to pileup files. Then, we generated depth chart files, which are tab-delimited text files recording the count of the number of nucleotides in columns and genomic positions in rows. We ran our variational inference algorithm on the depth chart files to identify SNVs.

5.2. **Performance on Synthetic DNA Data.**

5.2.1. *Comparison of Sensitivity and Specificity.* We compared the performance of our variational EM algorithm and the MCMC sampling algorithm in terms of sensitivity and specificity (Table 1). We used the posterior distribution test and χ^2 test to detect variants for a broad range of median read depths and NRAFs. The variational EM algorithm shows higher sensitivity and specificity than the MCMC algorithm in the events when NRAF is 0.1%. The variational EM algorithm has a higher specificity compared with the MCMC algorithm for a median read depth of $41,472\times$ at 0.3% NRAF and $55,489\times$ at 1.0% NRAF, but the sensitivity is slightly lower due to false negatives.

NRAF	Median Depth	Sensitivity		Specificity	
		MCMC	Variational	MCMC	Variational
0.1%	39	0.00	0.00	1.00	1.00
	408	0.00	0.07	1.00	1.00
	4129	0.14	0.29	1.00	1.00
	41449	0.86	1.00	0.97	1.00
0.3%	36	0.00	0.00	1.00	1.00
	410	0.00	0.00	1.00	1.00
	4156	1.00	1.00	0.99	0.98
	41472	1.00	0.93	0.85	0.91
1.0%	53	0.00	0.00	1.00	1.00
	535	0.21	0.29	1.00	1.00
	5584	1.00	1.00	0.98	0.98
	55489	1.00	0.93	0.87	0.95
10.0%	22	0.00	0.57	1.00	1.00
	260	1.00	1.00	1.00	1.00
	2718	1.00	1.00	1.00	1.00
	26959	1.00	1.00	1.00	1.00
100.0%	27	1.00	1.00	1.00	1.00
	298	1.00	1.00	1.00	1.00
	3089	1.00	1.00	1.00	1.00
	30590	1.00	1.00	1.00	1.00

TABLE 1. Sensitivity/Specificity comparison of variational EM algorithm with the MCMC algorithm on the synthetic DNA data set.

5.2.2. *Comparison of Approximated Posterior Distribution.* Figure 2 shows the approximate posterior distribution of the variational EM algorithm and samples of the MCMC algorithm. One variant position, 85, is taken as an example to show the comparison of the approximated posteriors. The variational EM and MCMC algorithms both identify all the variants when NRAF is 10.0% and 100.0%. The variational EM algorithm calls 90 false positive positions without a χ^2 test when NRAFs are 0.1% and 0.3% for low median read depth ($30\times$ and $400\times$). This is to be expected because it is highly unlikely to correctly identify a variant base with a population frequency of 1 in $1,000$ with less than a $1,000\times$ read depth.

A false positive, a non-mutated position that is called by the variational EM algorithm but not called by the MCMC algorithm, is shown in Figure 3. The variance of MCMC posterior estimate is higher than that of the variational posterior estimate. We tested 10 random initial values variational inference algorithm and found the approximate posterior distributions from the variational EM algorithm are essentially equivalent for all random initializations. It is notable that the shape of the proposed Beta variational distribution is well approximated by a Gaussian.

5.2.3. *Comparison to the State-of-the-Art Methods.* We compared the performance of our variational EM algorithm with the state-of-the-art variant detection methods, SAMtools [Li *et al.*, 2009], GATK [McKenna *et al.*, 2010], VarScan2 [Koboldt *et al.*, 2012], Strelka [Saunders *et al.*, 2012], MuTect [Cibulskis *et al.*, 2013], and RVD2 [He *et al.*, 2015], using synthetic DNA data set (Table 2). Among all of the methods compared, our variational EM algorithm has a higher sensitivity and specificity for a broad range of read depths and NRAFs. Our variational EM algorithm shows higher specificity than all the other tested methods at a very low NRAF (0.1%) level. However, our algorithm has a slightly lower specificity than the MCMC algorithm when the median read depth is $4,156\times$ at 0.3% NRAF, and a slightly lower sensitivity than the

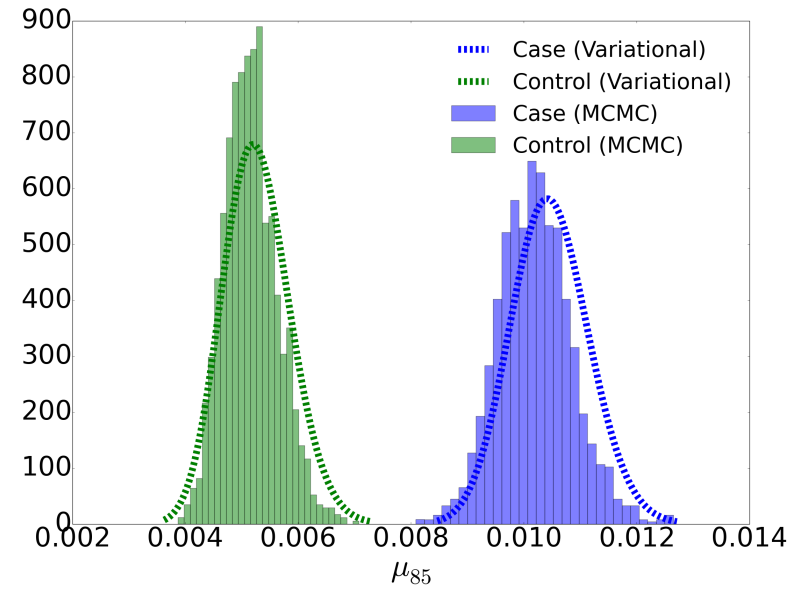

FIGURE 2. Approximated posterior distributions by the variational EM and MCMC algorithms on a true variant position 85 when the median read depth is $5,584\times$.

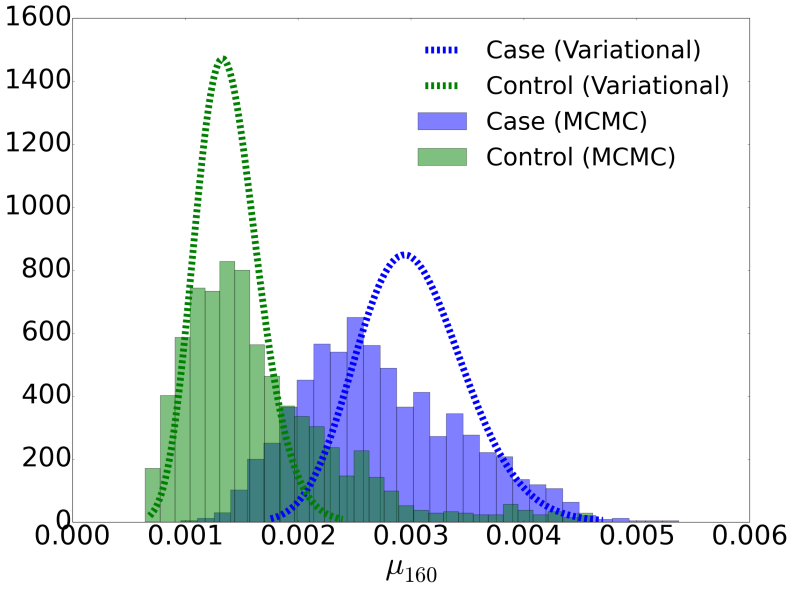

FIGURE 3. Approximated posterior distribution by the variational EM and MCMC algorithms on a false variant position 160 when the median read depth is $410\times$.

MCMC algorithm when the median read depth is $41,472\times$ at 0.3% NRAF and a median read depth of $55,489\times$ at 1.0% NRAF. The performance of other methods is stated in detail in He *et al.* [2015].

5.2.4. *Comparison of Timing.* The computational time for approximating the variational posterior distribution is increased by expanding the length of region and the median read depth (Figure 4). Our variational

NRAF	Median Depth	SAMtools	GATK	VarScan2 mpileup	VarScan2 somatic	Strelka	MuTect	MCMC RVD2	Variational RVD2
0.1%	39	0.00/1.00	0.00/1.00	0.00/1.00	0.00/1.00	0.00/1.00	0.00/0.99	0.00/1.00	0.00/1.00
	408	0.00/1.00	0.00/1.00	0.00/1.00	0.07/0.92	0.00/1.00	0.29/0.91	0.00/1.00	0.07/1.00
	4129	0.00/1.00	0.00/1.00	0.00/1.00	0.57/0.52	0.00/1.00	0.64/0.86	0.14/1.00	0.29/1.00
	41449	0.00/1.00	0.00/1.00	0.00/1.00	0.64/0.79	0.00/1.00	0.14/0.93	0.86/0.97	1.00/1.00
0.3%	36	0.00/1.00	0.00/1.00	0.00/1.00	0.00/1.00	0.00/1.00	0.43/1.00	0.00/1.00	0.00/1.00
	410	0.00/1.00	0.00/1.00	0.00/1.00	0.21/0.95	0.00/1.00	0.50/0.94	0.00/1.00	0.00/1.00
	4156	0.00/1.00	0.00/1.00	0.00/1.00	0.57/0.53	0.00/1.00	0.36/0.91	1.00/0.99	1.00/0.98
	41472	0.00/1.00	0.00/1.00	0.00/1.00	0.64/0.75	0.00/1.00	0.43/0.90	1.00/0.85	0.93/0.91
1.0%	53	0.00/1.00	0.00/1.00	0.00/1.00	0.00/0.99	0.00/1.00	0.29/0.98	0.00/1.00	0.00/1.00
	535	0.00/1.00	0.00/1.00	0.00/1.00	0.43/0.89	0.00/1.00	0.71/0.91	0.21/1.00	0.29/1.00
	5584	0.00/1.00	0.00/1.00	0.00/1.00	0.57/0.47	0.00/1.00	0.64/0.95	1.00/0.98	1.00/0.98
	55489	0.00/1.00	0.00/1.00	0.00/1.00	0.64/0.69	0.00/1.00	0.86/0.90	1.00/0.87	0.93/0.95
10.0%	22	0.21/1.00	0.43/1.00	0.00/1.00	0.36/1.00	0.29/1.00	0.86/0.99	0.00/1.00	0.57/1.00
	260	0.00/1.00	0.57/1.00	0.00/1.00	0.86/1.00	1.00/1.00	1.00/0.99	1.00/1.00	1.00/1.00
	2718	0.00/1.00	0.79/1.00	0.00/1.00	0.57/0.78	1.00/1.00	1.00/0.98	1.00/1.00	1.00/1.00
	26959	0.00/1.00	0.57/1.00	0.00/1.00	0.64/0.53	1.00/0.99	1.00/0.98	1.00/1.00	1.00/1.00
100.0%	27	1.00/0.99	1.00/1.00	1.00/1.00	1.00/1.00	1.00/1.00	1.00/0.98	1.00/1.00	1.00/1.00
	298	1.00/0.99	1.00/1.00	1.00/1.00	1.00/0.99	1.00/0.99	1.00/0.98	1.00/1.00	1.00/1.00
	3089	0.86/1.00	1.00/1.00	1.00/1.00	1.00/0.65	1.00/0.99	1.00/0.98	1.00/1.00	1.00/1.00
	30590	0.71/1.00	1.00/1.00	1.00/1.00	1.00/0.39	1.00/1.00	1.00/0.99	1.00/1.00	1.00/1.00

TABLE 2. Sensitivity/Specificity comparison of our variational RVD2 with other variant detection methods on synthetic DNA data.

EM algorithm is faster than the MCMC algorithm at the low median read depths of $27\times$ and $298\times$, and slower for the high median read depths of $3,089\times$ and $30,590\times$.

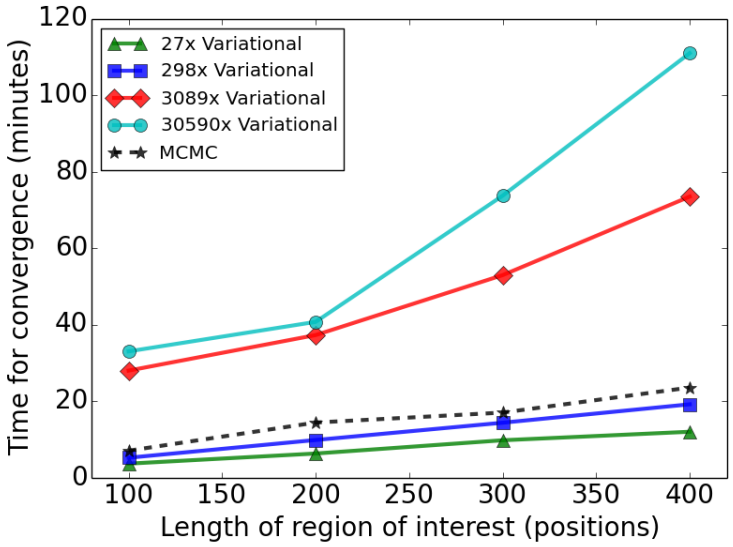

FIGURE 4. Timing for our variational EM algorithm and MCMC sampling algorithm. Sixty processers are used to estimate the model on the synthetic data set.

Table 3 shows the timing profile for each part of our variational EM algorithm when median read depth is $3,089\times$. Optimizing γ in the E-step (16) and optimizing M_j in the M-step (19) takes more than 95% of the time of one variational iteration in a test of a single processor, since the integration (12) is needed.

Computation resource	Region length	E-step			M-step				Total time (s)
		Optimize γ	Optimize δ	Update ELBO	Optimize μ_0	Optimize M_0	Optimize M	Update ELBO	
single processor	100	617.7 (63%)	4.232	10.42	0.264	0.159	332.8 (34%)	10.29	976.0
	200	1124 (65%)	8.936	18.64	0.418	0.256	570.0 (33%)	18.37	1741
	300	1728 (65%)	13.27	27.81	0.445	0.400	851.5 (32%)	27.65	2649
	400	2433 (66%)	17.99	38.55	0.737	0.635	1176 (32%)	38.17	3705
60 processors	100	29.93 (41%)	0.2470	11.67	0.3070	0.1890	19.56 (26%)	11.98	73.89
	200	44.69 (40%)	0.4170	22.14	0.5230	0.3040	24.04 (21%)	22.24	114.3
	300	63.47 (40%)	0.7160	33.31	0.5620	0.5040	29.41 (18%)	33.24	161.2
	400	94.66 (43%)	0.7270	42.78	0.8200	0.7060	40.04 (18%)	44.28	219.7

TABLE 3. Timing profile of 4 significant figures for one iteration of variational EM algorithm when median read depth is $3,089\times$. Single and multiple processors are both tested to estimate timing. Time for optimizing γ in the E-step and optimizing M in the M-step is highlighted in percentage.

5.3. **Variant Detection on the Longitudinal Directed Evolution Data.** We applied our variational EM algorithm to the MTH1 gene at Chr04:1,014,401-1,015,702 (1,302bp), which is the most frequently observed mutated gene by Kvitek and Sherlock [2013]. Our algorithm detected the same variants that were found by Kvitek and Sherlock [2013] (shown as highlighted in Table 4). Additionally, we detected 81 novel variants in 8 timepoints that the original publication did not detect. In Table 4, G7 is the baseline NRAF as the control sample when comparing with G70, G133, G266, G322, G385, and G448 in the respective hypotheses testing. The corresponding NRAFs of called variants at different time points are given by the estimate of the latent variable, $\hat{\mu_j} = E_q[\mu_j|r]$.

All of these variants, except the variant at position Chr04:1,014,740, decrease in NRAF following a maximum. The allele at position Chr04:1,014,740 is a beneficial variant that arises in NRAF to 99.603% at generation 448 within a constant glucose-limited environment. Moreover, we identified the first emergence of this beneficial variant as early as 0.522% in generation 133. Table 4 shows that we are able to detect many variants (NRAF $< 1.0\%$) early in the evolutionary time course. Furthermore, we detected a variant at position Chr04:1,015,666 at generation 70 as 0.091% NRAF, which is less than a 0.1% fraction.

We identified a pair of variants, Chr04:1,014,740 in gene MTH1 and Chr12:200,286 in gene ADE16, that increase in NRAF together in time (Figure 5). We hypotheses that the variants are concomitant in the same strain. In this pair of genes, gene MTH1 is a negative regulator of the glucose-sensing signal transduction pathway, and gene ADE16 is an enzyme of *de novo* purine biosynthesis. Glucose sensing induces gene expression changes to help yeast receive necessary nutrients, which could be a reason for this pair of genes to mutate together [Johnston, 1999]. Further experimental validation of this hypothesis would be required to definitively show that the mutations are concomitant.

The global precision parameter $\hat{M_0}$ could influence the estimate of μ_j due to its regularization effect. We show the influence of different $\hat{M_0}$ on variant position Chr04:1,014,740, $q(\mu_{1,014,740}|r)$ in Figure 6. We see that as we decrease the prior precision parameter $\hat{M_0}$, $\hat{\mu}_{1,014,740}$ increases as expected. But the effect of changing $\hat{M_0}$ over several orders of magnitude does not change μ_j greatly. Here $\hat{M_0} = 1.752$ in this dataset.

6. DISCUSSION

In this article, we propose a variational EM algorithm to estimate the non-reference allele frequency in the RVD2 model to identify rare nucleotide variants in heterogeneous pools. We derived a variational EM algorithm to avoid an intractable integration in our Bayesian model. Our results show that the variational

Index	Position	Ref	Alt	Allele frequency of generation (%)							
				G7	G70	G133	G196	G266	G322	G385	G448
1	1014407	T	C	0.127			0.620				
2	1014408	A	G	0.083	0.374						
3	1014422	C	G	0.160			1.036	7.431			
4	1014434	T	C	0.161						0.042	
5	1014436	A	G	0.161		0.127					
6	1014447	A	C	0.086			0.837	0.642			
7	1014455	C	T	0.044	0.345						
8	1014456	T	G	0.073	0.644	1.154					
9	1014457	G	T	0.031		0.777	2.675				
10	1014561	T	C	0.059	0.444						
11	1014580	A	G	0.161		0.096	0.168	0.481	0.088	0.038	0.081
12	1014582	T	C	0.161		0.149					
13	1014583	G	A	0.039	0.405	1.171	0.366				
14	1014595	A	G	0.065	0.096						
15	1014607	T	A	0.039	0.292	0.443					
16	1014615	G	T	0.079		0.865	1.327				
17	1014651	C	T	0.210			0.723				
18	1014689	T	A	0.161		0.101					
19	1014691	G	A	0.056		0.308					
20	1014698	C	T	0.160	0.13	0.629	0.565	0.767			
21	1014701	T	C	0.160		0.120					
22	1014707	A	C	0.160	0.825	1.461	7.434	6.920	0.176		
23	1014712	C	A	0.160	0.116	1.429	0.350	1.104			
24	1014740	G	C	0.158		0.522		23.849	98.286	99.715	99.603
25	1014741	C	T	0.028	0.997	1.622	0.780				
26	1014751	C	A	0.017		0.896					
27	1014765	A	C	0.159		0.431					
28	1014770	G	T	0.159	0.758	13.820	0.727	1.671			
29	1014791	C	T	0.014		0.918					
30	1014823	C	A	0.059		0.171					
31	1014856	T	C	0.160		0.119					
32	1014867	A	G	0.098	0.584	0.562	0.371	0.360			0.404
33	1014877	T	A	0.039			0.449				
34	1014920	G	C	0.049		0.324					
35	1014930	G	T	0.010		0.303					
36	1014958	C	A	0.160			1.884	1.036			
37	1014971	T	C	0.030	0.445						
38	1014968	A	T	0.037		0.316	3.446	1.462			
39	1014978	A	C	0.011		1.430					
40	1014997	G	A	0.008		0.368	2.287	2.564			
41	1015004	A	G	0.160	0.127						
42	1015036	G	A	0.026			0.640				
43	1015043	C	T	0.047		0.399					
44	1015051	G	A	0.161				0.084			
45	1015069	T	C	0.162						0.027	

Index	Position	Ref	Alt	Allele frequency of generation (%)							
				G7	G70	G133	G196	G266	G322	G385	G448
46	1015074	A	T	0.002		0.400	0.510				
47	1015077	G	T	0.047		1.06					
48	1015078	A	G	0.186			0.689				
49	1015086	T	C	0.055			0.525	0.437			0.307
50	1015092	A	C	0.003			0.424				
51	1015190	A	C	0.160		0.459	0.545				
52	1015220	G	A	0.161		0.061					
53	1015222	T	C	0.161					0.040		0.067
54	1015228	T	A	0.054			0.293				
55	1015236	A	C	0.095	0.935	2.211	1.425	1.100			
56	1015247	A	G	0.060			0.142				
57	1015276	T	C	0.109			0.519				
58	1015280	T	A	0.519			1.108				
59	1015284	G	A	0.159				0.190			
60	1015317	G	C	0.160			0.168	0.306			
61	1015321	G	A	0.059		0.803					
62	1015322	A	T	0.013		0.519	0.488				
63	1015360	C	A	0.161			0.067				
64	1015368	C	T	0.064		0.147					
65	1015370	T	C	0.161		0.149					
66	1015371	G	C	0.062				0.147			
67	1015386	G	T	0.047	0.375	3.012	2.098	1.022			
68	1015411	A	G	0.042	0.214						
69	1015423	G	A	0.161			0.159				
70	1015424	T	C	0.050						0.194	
71	1015434	A	T	0.161		0.151					
72	1015477	C	A	0.043	0.440						
73	1015478	C	T	0.161			0.162				
74	1015512	G	T	0.161						0.006	
75	1015519	T	C	0.161		0.080					
76	1015521	G	A	0.018	0.371						
77	1015522	A	G	0.040			0.343				
78	1015539	G	C	0.040		0.618					
79	1015555	G	A	0.160		0.110					
80	1015556	A	G	0.161			0.117				
81	1015563	G	A	0.041			0.247				
82	1015623	T	G	0.034							0.194
83	1015627	T	C	0.161					0.060		
84	1015657	G	T	0.032		0.740	0.583	0.461			
85	1015666	G	A	0.065	0.091		0.105				
86	1015681	C	T	0.161					0.030		
87	1015691	T	C	0.037		0.358					
88	1015699	A	G	0.161			0.068				
89	1015700	C	A	0.015	0.949	1.965	1.044	0.379			
90	1015701	A	G	0.161		0.068					
Num of variants				0	22	24	42	21	6	6	6

TABLE 4. Identified variants and corresponding NRAFs in gene MTH1 on Chromosome 4. A blank cell indicates that the position of that time point is not called significantly different than G7. The positions highlighted as blue were also identified by Kvitek, 2013. The other 81 positions are novel identified variants in 8 timepoints.

EM algorithm (i) is able to identify rare variants at a 0.1% NRAF level with comparable sensitivity and specificity to a MCMC sampling algorithm; (ii) has a higher specificity in comparison with many state-of-the-art algorithms in a broad range of NRAFs; and (iii) detects SNVs early in the evolutionary time course, as well as tracks NRAF in a real longitudinal yeast data set.

We have chosen parametric forms for the variational distributions. This choice has left us with a complex integral in our variational optimization problem. In future work, we plan to explore other approximations of the variational distributions that render the integral easier to compute. One could use cubic splines to numerically approximate the function and then integrate that surrogate [McKinley and Levine, 1998]. Another strategy is to consider a Laplace approximation for the variational distribution, as we and others have done previously [Saddiki *et al.*, 2014; Wang and Blei, 2013].

Improving the speed of the estimating algorithm enables us to interrogate whole-genome sequencing data. By doing this, we hope to reveal the dynamics of arising variants at the genome-wide scale to show the

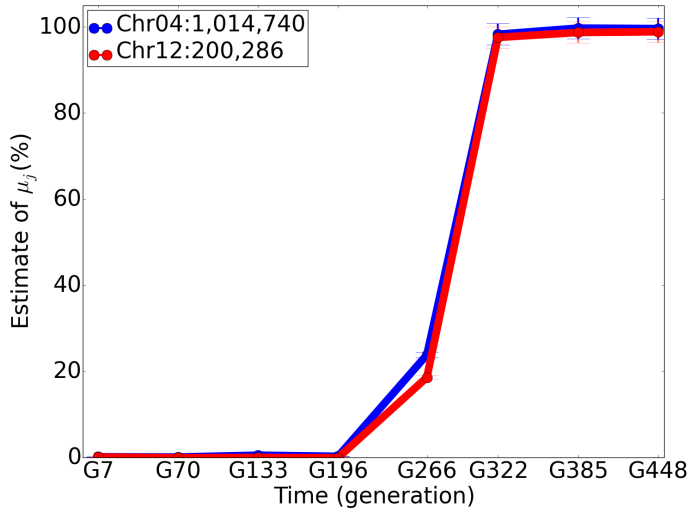

FIGURE 5. The NRAF trend of concomitant variants in gene MTH1 and ADE16. The 95% Bayesian credible intervals are shown.

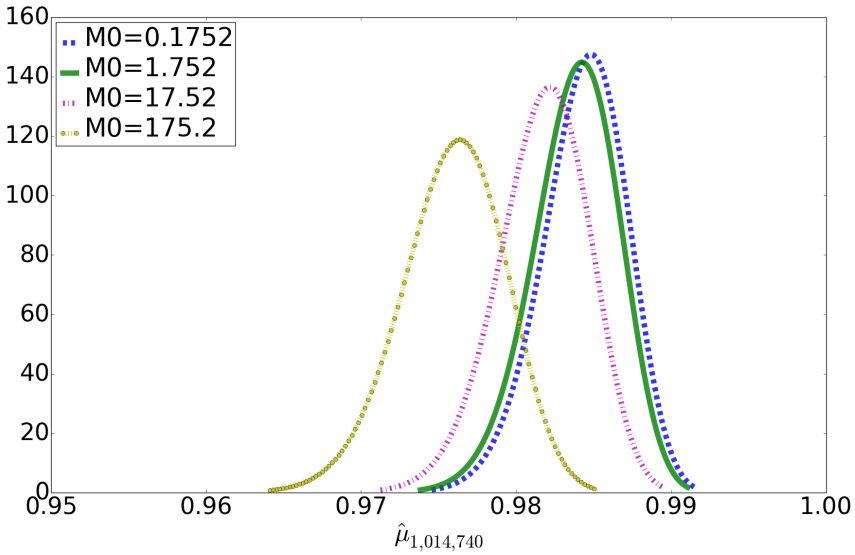

FIGURE 6. Influence of M_0 on the estimate of μ_j. Posterior distributions of the variant at position Chr04:1,014,740, $\hat{\mu}_{1,014,740}$, with different $\hat{M}_0$ are shown.

genetic basis of clonal interference. Our method can also be extended to study drug resistance by characterizing tumor heterogeneity in targeted anti-cancer chemotherapy samples. We will be able to find the causative variants that lead to drug resistance and understand the causes of resistance at the single nucleotide level.

7. ACKNOWLEDGMENTS

F.Z. and P.F. were supported by PhRMA Foundation Informatics Grant 2013080079. We would like to thank Dan Kvitek for kindly sharing the data set used for the longitudinal yeast evolution analysis. We acknowledge Chuangqi Wang and Ross Lagoy for valuable comments on the manuscript.

APPENDIX A. DERIVATION OF THE COMPONENTS IN ELBO

Here we derive the requisite expectations with respect to the variational distribution for the ELBO. Recall that $\theta_{ji} \sim \text{Beta}(\mu_j, M_j)$. Then we can write,

$$
\begin{aligned}
E_q\left[\log p\left(\theta|\mu; M\right)\right] &= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log p\left(\theta_{ji}|\mu_j; M_j\right)\right] \\
&= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\theta_{ji}^{M_j\mu_j-1}(1-\theta_{ji})^{M_j(1-\mu_j)-1}\right)\right] \\
&= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right] \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\theta_{ji}^{M_j\mu_j-1}(1-\theta_{ji})^{M_j(1-\mu_j)-1}\right)\right] \\
&= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right] \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{E_q\left[(M_j\mu_j-1)\log\theta_{ji}\right] + E_q\left[(M_j(1-\mu_j)-1)\log(1-\theta_{ji})\right]\right\} \\
&= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right] \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{M_j E_q\left[\mu_j\right] E_q\left[\log\theta_{ji}\right] - E_q\left[\log\theta_{ji}\right]\right\} \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{\left(M_j - 1 - M_j E_q\left[\mu_j\right]\right) E_q\left[\log(1-\theta_{ji})\right]\right\} \\
&= N * \sum_{j=1}^{J} E_q\left[\log\left(\frac{\Gamma(M_j)}{\Gamma(\mu_j M_j)\Gamma(M_j(1-\mu_j))}\right)\right] \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{M_j E_q\left[\mu_j\right] E_q\left[\log\theta_{ji}\right] - E_q\left[\log\theta_{ji}\right]\right\} \\
&\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{\left(M_j - 1 - M_j E_q\left[\mu_j\right]\right) E_q\left[\log(1-\theta_{ji})\right]\right\}
\end{aligned}
\tag{26}
$$

Since $\mu_j \sim \text{Beta}(\mu_0, M_0)$,

$$\begin{aligned} E_q\left[\log p\left(\mu;\mu_0,M_0\right)\right] &= \sum_{j=1}^{J} E_q\left[\log p\left(\mu_j;\mu_0,M_0\right)\right] \\ &= \sum_{j=1}^{J} E_q\left[\log\left(\frac{\Gamma(M_0)}{\Gamma(\mu_0 M_0)\Gamma(M_0(1-\mu_0))}\mu_j^{M_0\mu_0-1}(1-\mu_j)^{M_0(1-\mu_0)-1}\right)\right] \\ &= J * \log\frac{\Gamma(M_0)}{\Gamma(\mu_0 M_0)\Gamma(M_0(1-\mu_0))} \\ &\quad + \sum_{j=1}^{J}\left\{(M_0\mu_0-1)E_q\left[\log\mu_j\right] + (M_0(1-\mu_0)-1)E_q\left[\log(1-\mu_j)\right]\right\}. \end{aligned} \tag{27}$$

Finally, since $r_{ji}|n_{ji} \sim \text{Binomial}(\theta_{ji}, n_{ji})$,

$$\begin{aligned} E_q\left[\log p\left(r|\theta,n\right)\right] &= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log p\left(r_{ji}|\theta_{ji},n_{ji}\right)\right] \\ &= \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[\log\left(\frac{\Gamma(n_{ji}+1)}{\Gamma(r_{ji}+1)\Gamma(n_{ji}-r_{ji}+1)}\theta_{ji}^{r_{ji}}(1-\theta_{ji})^{n_{ji}-r_{ji}}\right)\right] \\ &= \sum_{j=1}^{J}\sum_{i=1}^{N}\log\left(\frac{\Gamma(n_{ji}+1)}{\Gamma(r_{ji}+1)\Gamma(n_{ji}-r_{ji}+1)}\right) \\ &\quad + \sum_{j=1}^{J}\sum_{i=1}^{N} E_q\left[r_{ji}\log\theta_{ji} + (n_{ji}-r_{ji})\log(1-\theta_{ji})\right] \\ &= \sum_{j=1}^{J}\sum_{i=1}^{N}\log\left(\frac{\Gamma(n_{ji}+1)}{\Gamma(r_{ji}+1)\Gamma(n_{ji}-r_{ji}+1)}\right) \\ &\quad + \sum_{j=1}^{J}\sum_{i=1}^{N}\left\{r_{ji}E_q\left[\log\theta_{ji}\right] + (n_{ji}-r_{ji})E_q\left[\log(1-\theta_{ji})\right]\right\}. \end{aligned} \tag{28}$$

REFERENCES

Vikas Bansal. A statistical method for the detection of variants from next-generation resequencing of dna pools. *Bioinformatics*, 26(12):i318–i324, 2010.

Kristian Cibulskis, Michael S Lawrence, Scott L Carter, Andrey Sivachenko, David Jaffe, Carrie Sougnez, Stacey Gabriel, Matthew Meyerson, Eric S Lander, and Gad Getz. Sensitive detection of somatic point mutations in impure and heterogeneous cancer samples. *Nature biotechnology*, 31(3):213–219, 2013.

Bradley Efron. *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press, 2010.

Patrick Flaherty, Georges Natsoulis, Omkar Muralidharan, Mark Winters, Jason Buenrostro, John Bell, Sheldon Brown, Mark Holodniy, Nancy Zhang, and Hanlee P Ji. Ultrasensitive detection of rare mutations using next-generation targeted resequencing. *Nucleic acids research*, page gkr861, 2011.

Elodie Ghedin, Jennifer Laplante, Jay DePasse, David E Wentworth, Roberto P Santos, Martha L Lepow, Joanne Porter, Kathleen Stellrecht, Xudong Lin, Darwin Operario, et al. Deep sequencing reveals mixed

infection with 2009 pandemic influenza a (h1n1) virus strains and the emergence of oseltamivir resistance. *Journal of Infectious Diseases*, 203(2):168–174, 2011.

Yuting He, Fan Zhang, and Patrick Flaherty. Rvd2: An ultra-sensitive variant detection model for low-depth heterogeneous next-generation sequencing data. *Bioinformatics*, page btv275, 2015.

Mark Johnston. Feasting, fasting and fermenting: glucose sensing in yeast and other cells. *Trends in Genetics*, 15(1):29–33, 1999.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

Katy C Kao and Gavin Sherlock. Molecular characterization of clonal interference during adaptive evolution in asexual populations of saccharomyces cerevisiae. *Nature genetics*, 40(12):1499–1504, 2008.

Daniel C Koboldt, Qunyuan Zhang, David E Larson, Dong Shen, Michael D McLellan, Ling Lin, Christopher A Miller, Elaine R Mardis, Li Ding, and Richard K Wilson. Varscan 2: somatic mutation and copy number alteration discovery in cancer by exome sequencing. *Genome research*, 22(3):568–576, 2012.

Daniel C Koboldt, Karyn Meltz Steinberg, David E Larson, Richard K Wilson, and Elaine R Mardis. The next-generation sequencing revolution and its impact on genomics. *Cell*, 155(1):27–38, 2013.

Dieter Kraft et al. *A software package for sequential quadratic programming*. DFVLR Obersfaffeuhofen, Germany, 1988.

Daniel J Kvitek and Gavin Sherlock. Whole genome, whole population sequencing reveals that loss of signaling networks is the major adaptive strategy in a constant environment. *PLoS Genet*, 9(11):e1003972, 2013.

David E Larson, Christopher C Harris, Ken Chen, Daniel C Koboldt, Travis E Abbott, David J Dooling, Timothy J Ley, Elaine R Mardis, Richard K Wilson, and Li Ding. Somaticsniper: identification of somatic point mutations in whole genome sequencing data. *Bioinformatics*, 28(3):311–317, 2012.

Heng Li, Bob Handsaker, Alec Wysoker, Tim Fennell, Jue Ruan, Nils Homer, Gabor Marth, Goncalo Abecasis, Richard Durbin, et al. The sequence alignment/map format and samtools. *Bioinformatics*, 25(16):2078–2079, 2009.

Aaron McKenna, Matthew Hanna, Eric Banks, Andrey Sivachenko, Kristian Cibulskis, Andrew Kernytsky, Kiran Garimella, David Altshuler, Stacey Gabriel, Mark Daly, et al. The genome analysis toolkit: a mapreduce framework for analyzing next-generation dna sequencing data. *Genome research*, 20(9):1297–1303, 2010.

Sky McKinley and Megan Levine. Cubic spline interpolation. *College of the Redwoods*, 45(1):1049–1060, 1998.

Nicholas Navin, Alexander Krasnitz, Linda Rodgers, Kerry Cook, Jennifer Meth, Jude Kendall, Michael Riggs, Yvonne Eberling, Jennifer Troge, Vladimir Grubor, et al. Inferring tumor progression from genomic heterogeneity. *Genome research*, 20(1):68–80, 2010.

Carsten Peterson and Eric Hartman. Explorations of the mean field theory learning algorithm. *Neural Networks*, 2(6):475–494, 1989.

Andrew Roth, Jiarui Ding, Ryan Morin, Anamaria Crisan, Gavin Ha, Ryan Giuliany, Ali Bashashati, Martin Hirst, Gulisa Turashvili, Arusha Oloumi, et al. Jointsnvmix: a probabilistic model for accurate detection of somatic mutations in normal/tumour paired next-generation sequencing data. *Bioinformatics*, 28(7):907–913, 2012.

Hachem Saddiki, Jon McAuliffe, and Patrick Flaherty. Glad: a mixed-membership model for heterogeneous tumor subtype classification. *Bioinformatics*, page btu618, 2014.

Christopher T Saunders, Wendy SW Wong, Sajani Swamy, Jennifer Becq, Lisa J Murray, and R Keira Cheetham. Strelka: accurate somatic small-variant calling from sequenced tumor–normal sample pairs. *Bioinformatics*, 28(14):1811–1817, 2012.

Jay Shendure and Hanlee Ji. Next-generation dna sequencing. *Nature biotechnology*, 26(10):1135–1145, 2008.

Yuichi Shiraishi, Yusuke Sato, Kenichi Chiba, Yusuke Okuno, Yasunobu Nagata, Kenichi Yoshida, Norio Shiba, Yasuhide Hayashi, Haruki Kume, Yukio Homma, et al. An empirical bayesian framework for somatic mutation detection from cancer genome sequencing data. *Nucleic acids research*, 41(7):e89–e89, 2013.

Chong Wang and David M Blei. Variational inference in nonconjugate models. *The Journal of Machine Learning Research*, 14(1):1005–1031, 2013.

Qingguo Wang, Peilin Jia, Fei Li, Haiquan Chen, Hongbin Ji, Donald Hucks, Kimberly Brown Dahlman, William Pao, Zhongming Zhao, et al. Detecting somatic point mutations in cancer genome sequencing data: a comparison of mutation callers. *Genome Med*, 5(10):91, 2013.

Zhi Wei, Wei Wang, Pingzhao Hu, Gholson J Lyon, and Hakon Hakonarson. Snver: a statistical tool for variant calling in analysis of pooled or individual next-generation sequencing data. *Nucleic acids research*, 39(19):e132–e132, 2011.